arXiv:1405.2856v1 [cs.DL] 12 May 2014

Mapping the UK Webspace: Fifteen Years of British Universities on the Web

Scott A. Hale
Oxford Internet Institute
University of Oxford
1 St Giles, Oxford UK
scott.hale@oii.ox.ac.uk

Taha Yasseri
Oxford Internet Institute
University of Oxford
1 St Giles, Oxford UK
taha.yasseri@...

Josh Cowls
Oxford Internet Institute
University of Oxford
1 St Giles, Oxford UK
josh.cowls@...

Eric T. Meyer
Oxford Internet Institute
University of Oxford
1 St Giles, Oxford UK
eric.meyer@...

Ralph Schroeder
Oxford Internet Institute
University of Oxford
1 St Giles, Oxford UK
ralph.schroeder@...

Helen Margetts
Oxford Internet Institute
University of Oxford
1 St Giles, Oxford UK
helen.margetts@...

ABSTRACT

This paper maps the national UK web presence on the basis of an analysis of the `.uk` domain from 1996 to 2010. It reviews previous attempts to use web archives to understand national web domains and describes the dataset. Next, it presents an analysis of the `.uk` domain, including the overall number of links in the archive and changes in the link density of different second-level domains over time. We then explore changes over time within a particular second-level domain, the academic subdomain `.ac.uk`, and compare linking practices with variables, including institutional affiliation, league table ranking, and geographic location. We do not detect institutional affiliation affecting linking practices and find only partial evidence of league table ranking affecting network centrality, but find a clear inverse relationship between the density of links and the geographical distance between universities. This echoes prior findings regarding offline academic activity, which allows us to argue that real-world factors like geography continue to shape academic relationships even in the Internet age. We conclude with directions for future uses of web archive resources in this emerging area of research.

Categories and Subject Descriptors

H.5.4 [**Information Interfaces and Presentation (e.g. HCI)**]: Hypertext/Hypermedia; H.5.3 [**Information Systems**]: Group and Organization Interfaces—Web-based interaction

General Terms

Human Factors, Measurement

© Scott A. Hale, Taha Yasseri, Josh Cowls, Eric T. Meyer, Ralph Schroeder, and Helen Margetts, 2014. This is the authors' version of the work. It is posted here for your personal use. Not for redistribution. The definitive version is published in WebSci '14, `http://dx.doi.org/10.1145/2615569.2615691`.

Keywords

Web Archives; World Wide Web; Network Analysis; Hyperlink Analysis; Big Data; Academic Web

1. INTRODUCTION

The World Wide Web is enormous and is in constant flux, with more web content lost to time than is currently accessible via the live web. The growing body of archived web material available to researchers is thus potentially immensely valuable as a record of important aspects of modern society, but there have previously been few tools available to facilitate research using archived web materials [6]. Nevertheless, with the development of new tools and techniques such as those used in this paper, the use of web archives both to understand the history of the web itself as well as to shed light on broader changes in society is emerging as a promising research area [7]. The web is likely to provide insight into social changes just as other historical artifacts, such as newspapers and books, have done for scholars interested in the pre-digital world. As the web becomes increasingly embedded in all spheres of everyday life and the number of webpages continues to grow, there is a compelling case to be made for examining changes in both the structure and content of the web. However, while interfaces such as the Wayback Machine[1] allow access to individual webpages one at a time, there have been relatively few attempts to work with large collections of web archive data using computational approaches across the corpus. This paper provides a longitudinal analysis of the UK national web domain, `.uk`, and the academic second-level domain, `.ac.uk`, in order to show the benefits and challenges of this type of analysis.

2. BACKGROUND

2.1 Archiving national web domains

National web domains represent one approach to web archive analysis for researchers seeking an overview of a single country's web presence. A particular national web domain offers the potential of both diversity and completeness in its coverage [2], although there are limitations in terms of gener-

[1] `http://web.archive.org/`

alizability beyond the country in question and frequently in terms of the completeness of the analysis based on technical factors (see below). At the same time, however, limiting the focus to a single country also has the potential to introduce fewer contextual differences (such as language, Internet penetration rates, broadband penetration rates, political openness, economic differences, and so forth), and thus is a sound strategy for demonstrating the potential of this type of analysis, which has not previously been done.

Research in this area is at an early stage, and there are conceptual challenges associated with analyzing national web domains. The content and structure of country-code top-level domains (ccTLDs) such as `.uk` for the United Kingdom and `.fr` for France are governed more by traditions than rules [14], complicating efforts to reach a comprehensive definition of what they represent. Brügger [5] discusses the difficulty, for example, of deciding how national presences should be delimited. In the case presented here, the domain name `.uk` is used, but this does not cover all the webpages originating in the United Kingdom as several British companies, organizations, and individuals operate domains in generic top-level domains (`.com`, `.org`, etc.) or elsewhere. Moreover webpages ending with `.uk` are also used for websites which arguably belong to a different country, as when multinational companies headquartered outside the UK have affiliates within the UK with a `.uk` address. Finally, it might be contended that not only webpages with a `.uk` address be examined, but also those that link to and from these webpages. However, for the purposes of this research, these limitations can mostly be noted for future research and do not seriously limit the ability to understand the broad patterns within the UK national web presence.

Another issue that must be decided when undertaking analysis of web domains is the appropriate level of detail. This includes the temporal resolution to use for analysis (since while the web is constantly changing, the number of snapshots available in Internet Archive data vary over time based on the crawl settings in place when the data were gathered) and what level of detail to extract from webpages (i.e., determining the appropriate level of resolution of page content, link information, page metadata, and so forth). Previous research on the `.uk` ccTLD has examined monthly snapshots over a one year period finding page-level hyperlinks change frequently month to month [3]. As Brügger [4] notes, there are several reasons why archived websites are different from other archived material in respect to these details: choices must be made not just about what to capture but also technical issues about what can be archived and how the archiving process itself shapes the later availability of the archived materials.

2.2 UK web domain

For the `.uk` domain that will be examined here, the source of the data is the archive files of the UK domain that were obtained from the Internet Archive by the British Library with the specific purpose of creating the basis of a national archive of the web in the UK. This data is currently being expanded via ongoing web archiving activities being performed by the British Library under the terms of the 2003 UK legal deposit law,[2] which was implemented via new regulations that went into effect in April 2013.[3]

The `.uk` country-code top-level domain is managed by the Internet registrar Nominet.[4] Below the `.uk` top-level domain are several second-level domains (SLDs), the largest of which are `.co.uk` (commercial enterprises), `.org.uk` (non-commercial organizations), `.gov.uk` (government bodies), and `.ac.uk` (academic establishments).[5] This paper examines the data aggregated to the level of third-level domains such as `nominet.org.uk` (Nominet), `fco.gov.uk` (the Foreign and Commonwealth Office of the UK government), or `ox.ac.uk` (the University of Oxford).

In the case of web archives (or indeed of other archived material which takes the approach of archiving all that can be archived, without a particular topic in mind), it is not scholarly interest that sets the agenda, but rather the goal of the archiving institution. This means that the scope of the archived material and the level of detail available, as with other historical materials, is a function of the archiving processes used to gather and store the data. Thus, unlike web archive research done on the live web using researcher-implemented data collection mechanisms [e.g., 8, 9], for the purpose of this study the dataset itself should be seen as a given. However, it can be mentioned that the Internet Archive's data comprises the most comprehensive archive of the web available [1].

3. DATA

3.1 Data preparation

The data for this study originally comes from the Internet Archive, which began crawling pages from all domains in 1996 [12]. Copies of the approximately 30 terabytes of compressed data relating to the `.uk` country-code top-level domain (ccTLD) was provided to the British Library and forms the "JISC UK Web Domain Dataset."[6]

Hale et al. [11] cleaned the data by removing error pages (e.g. 404 Not Found pages) as well as pages not within the `.uk` ccTLD. They produced a plain-text list of all page urls remaining in the collection and the date and times they were crawled, and an additional plain-text list of all outgoing hyperlinks starting from pages within the dataset.

For this study, we started with this list of hyperlinks and filtered it to only include links between different third-level domains. We further grouped pages crawled at similar times (within 1,000 seconds) together and assigned the hyperlink pair a weight based on the number of hyperlinks between the two third-level domains in that time period. For each year, we take the crawl with the largest number of captured hyperlinks between any two domains. We also formed one list of all third-level domains present in the dataset each year and the number of pages crawled within each third-level domain. We loaded these lists into Apache Hive for further analysis.

[2]Legal Deposit Libraries Act 2003, `http://www.legislation.gov.uk/ukpga/2003/28/contents`

[3]The Legal Deposit Libraries (Non-Print Works) Regulations, `http://www.legislation.gov.uk/uksi/2013/777/contents/made`

[4]`http://www.nominet.org.uk/`

[5]`http://www.nominet.org.uk/uk-domain-names/about-domain-names/uk-domain-subdomains/second-level-domains`

[6]`http://data.webarchive.org.uk/opendata/ukwa.ds.2/`

3.2 Data analysis

In what follows, we undertake a longitudinal network analysis, charting the .uk domain and its core second-level domains over time. As Brügger [4] points out, this type of analysis is not concerned with who produced what, nor with how the web content was used, but rather with what was created and thus "the web which is"—or rather was—"actually available to users."

First, we present an overall longitudinal view of the second-level domains within the .uk domain. We investigate the growth of the entire domain between 1996 and 2010, broken down into its four largest constituent parts, .co.uk, .org.uk, .gov.uk, and .ac.uk. Analysis of these SLDs allows us to investigate the role of different sectors of British society in the growth of the UK web presence.

The second section looks at the linking practices between these SLDs. It asks about the internal link density of each SLD, and analyses how they interact with each other: whether, for example, there are more links between certain subdomains, and whether linking is reciprocal between domains or imbalanced.

The third and final section of the findings takes a closer look at the SLD .ac.uk. It builds on earlier longitudinal analyses of academic webpages, which have investigated, for example, the stability of outlinks [18, 20]. Our findings update earlier studies by extending the period of analysis to the end of 2010 and assessing the effects of new variables, including institutional affiliation, league table ranking, and geographic location on link practices between different universities.

4. RESULTS

4.1 Overview of .uk

Figure 1 displays the overall growth of the .uk ccTLD, showing the total number of nodes (on a logarithmic scale) within each SLD from 1996 to 2010. It also shows the size of the entire .uk domain space (on a linear scale). There is a clear change in the trend of the growth around 2001 for .co.uk and .org.uk. Furthermore, .ac.uk and .gov.uk seem to almost stabilize in size at around the same time.

Figure 2 shows the relative size of the second-level domains .co.uk, .org.uk, .ac.uk, and .gov.uk across the fifteen year period, standardized as each SLD's proportion of the total nodes (i.e., domains/websites, not webpages) in the collection in each year. While these are not the only second-level domains in use within the .uk domain, they are the four largest in terms of number of nodes across the whole period.

As Figure 2 shows, .co.uk is the predominant second-level domain throughout the entire period, with .co.uk sites never accounting for less than 85% of the total. However, also apparent is the large proportion of governmental and, especially, academic sites in the early recorded history of the UK web. This is consistent with the role that universities played in the early establishment, adoption, and development of the web [13]. Over time, however, this early presence was greatly overshadowed in terms of absolute numbers of nodes when compared to the continued growth of the .co.uk and .org.uk domains.

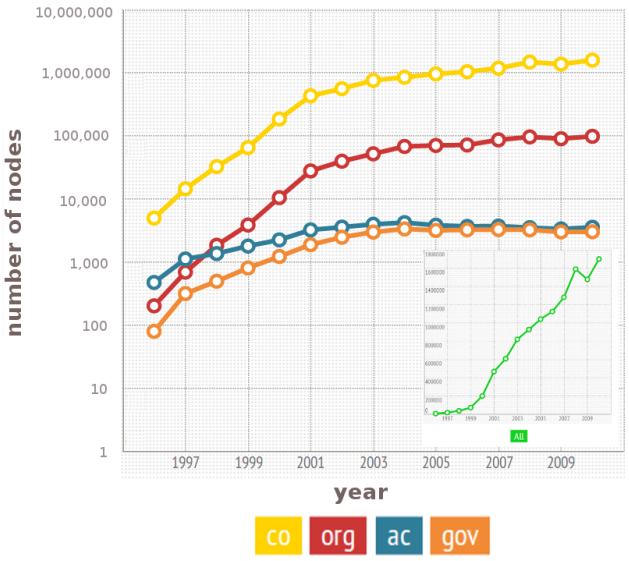

Figure 1: Number of nodes (third-level domains) within each second-level domain over time. The inset shows the sum over all second-level domains.

4.2 Link density among and between second-level domains

Up to this point the analysis has drawn only on node data; that is, the number of websites making up each domain. However, link analysis can offer insight into how well integrated each SLD is with itself and with other domains. A link from one site to another has been used as an indicator of awareness between blogs [10] and recognition between academic sites [20]. Figure 3 shows, for each subdomain, how many links there are for every node over time, where a fluctuating relationship between the number of nodes and links to other nodes for each second-level domain is visible. Over the whole period, the .ac.uk academic SLD and, from 1997 onwards, the .gov.uk governmental SLD are the most internally dense SLDs. This observation may reflect the fact that registration for the .ac.uk and .gov.uk subdo-

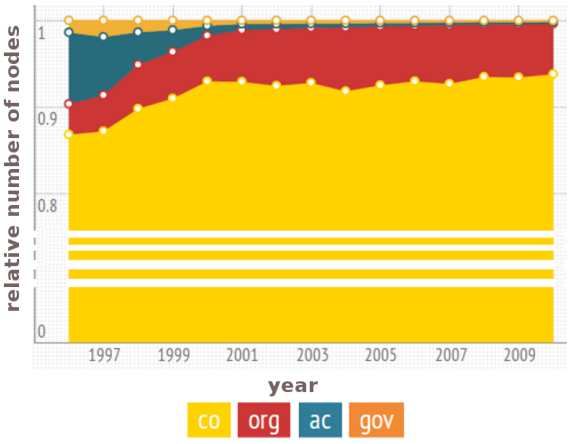

Figure 2: Relative size of second-level domains in the .uk top-level domain over time.

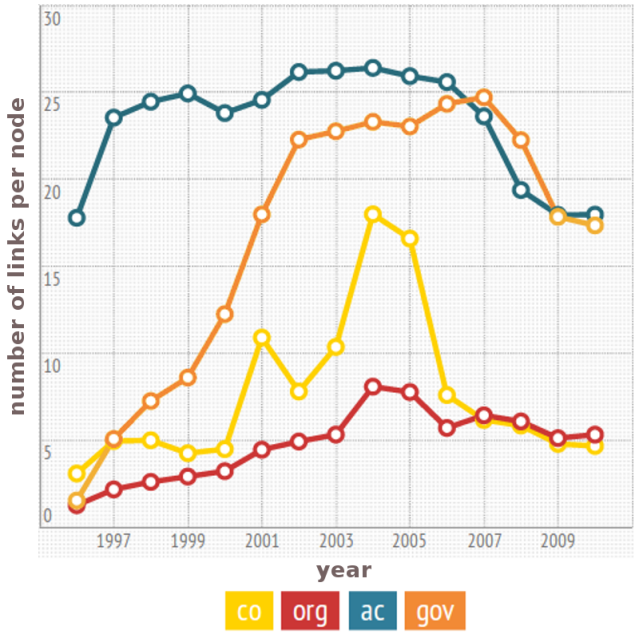

Figure 3: Number of within-SLD links per node in four .uk SLDs, 1996–2010.

mains is restricted, whereas .org.uk and .co.uk sites can be registered easily by any party. In addition, the .ac.uk and .gov.uk subdomains are likely constituted by a narrower and more cohesive set of institutions, creating, on average, a stronger basis for linking within the SLDs. Furthermore, there is likely more competition and thus less reason to link within the .co.uk commercial subdomain compared to .ac.uk or .gov.uk. Higher link density within the .org and .gov domains in comparison to the .com domain has previously been observed in a smaller scale, topical study about climate change [19].

Also of note is the general rise of links in the middle of the period, particularly in the substantial .co.uk subdomain. This peaks sharply in 2004 before falling sharply back to around pre-2001 levels by 2009. This trend has no easy explanation, suggesting that further research is required to explain this pattern. Possible explanations include the norm of including lists of links on webpages such as blogs fell out of favor in the middle of this period or that more websites increasingly linked outside of the .uk ccTLD.

Not only can web domain data tell us how well integrated an SLD is internally, but we can also investigate how well SLDs are connected to each other. Figures 4a and 4b show the quantity of links between SLDs for 2010, the last year in the dataset, where the size of an arc relates to the volume of links from one SLD to another. The color of each arc relates to links sent in one direction, from the host SLD outwards. For example green arcs show links from the .co.uk domain to others. Figure 4a shows the absolute volume of links, while the size of the arcs in Figure 4b are normalized in relation to the number of nodes in the target subdomain. (Note that Figure 4a does not display links within a single SLD, as the volume of links between .co.uk sites dwarfs all other relationships. As Figure 4b controls for the number of nodes in each SLD, the adjusted .co.uk arc is much smaller and links within a single SLD are therefore included.)

Figure 4a shows that the largest volume of links between SLDs in 2010 flowed from .co.uk sites to .org.uk sites, and this relationship is fairly reciprocal, with .org.uk sites sending almost as many links back. Links between other domains are much lower in terms of absolute volume. When controlling for the size of the target subdomain, however, the picture changes somewhat. As Figure 2 above showed, by 2010 the number of nodes in the .org.uk subdomain far outweighed those in the .ac.uk and .gov.uk subdomains. Figure 4b, adjusting for this, shows that the .gov.uk and, to a lesser extent, the .ac.uk subdomains punch above their weight, receiving proportionally more links from .co.uk and .org.uk sites. Once again, the more restrictive registration policies for these SLDs may be a factor here, driving up the average quality and 'linkworthiness' of sites in these subdomains as compared to .co.uk and .org.uk sites although it may also be related to other factors such as the comparative homogeneity of these SLDs, the perception of objectivity or balance on academic or government websites as opposed to sites oriented towards sales or persuasion, or even the international standing of many UK universities, although understanding these factors would require further investigation.

For the .gov.uk subdomain, the finding that sites link out less than they are linked to suggests a lack of 'outward-lookingness,' compared to the other sectors. In contrast, Escher et al. [8] found the UK Foreign and Commonwealth Office to be relatively more outward-looking than its equivalents in Australia and the US. However, foreign offices, with their outward facing role could easily be an exception to more general government-wide propensity not to link out.

In addition, it is worth noting the relatively heavy proportion of links within the .ac.uk SLD shown in Figure 4b. This propensity of academic institutions to link heavily to other academic institutions (more so than the other domains) reflects (taking a positive view) a strong network among academic institutions, but also potentially (taking a negative view) a tendency towards inward-looking, within-domain links. We examine these links in more depth in the next section.

4.3 The case of the academic subdomain

At this stage we turn our attention to one particular subdomain, the .ac.uk academic subdomain of the UK web. To be eligible for a third-level domain within .ac.uk, an organization must have a permanent physical presence in the UK and either have the majority of its activities publicly funded by UK government funding bodies or be a Learned Society. In addition, the organization must satisfy at least one of the following criteria: the organization must provide tertiary-level education with central government funding, conduct publicly funded academic research, have a primary purpose of supporting tertiary-level educational establishments, or have the status of a Learned Society ("a society that exists to promote an academic discipline or group of disciplines").[7]

The academy was at the forefront of the development of the web, and, as Figure 2 shows, .ac.uk sites constituted a sizeable minority of .uk sites in 1996. Over time, this proportion waned, even as more British universities established a substantial web presence. In this subsection we use the longitudinal data collected to examine the relationship be-

[7] https://community.ja.net/library/janet-services-documentation/eligibility-guidelines

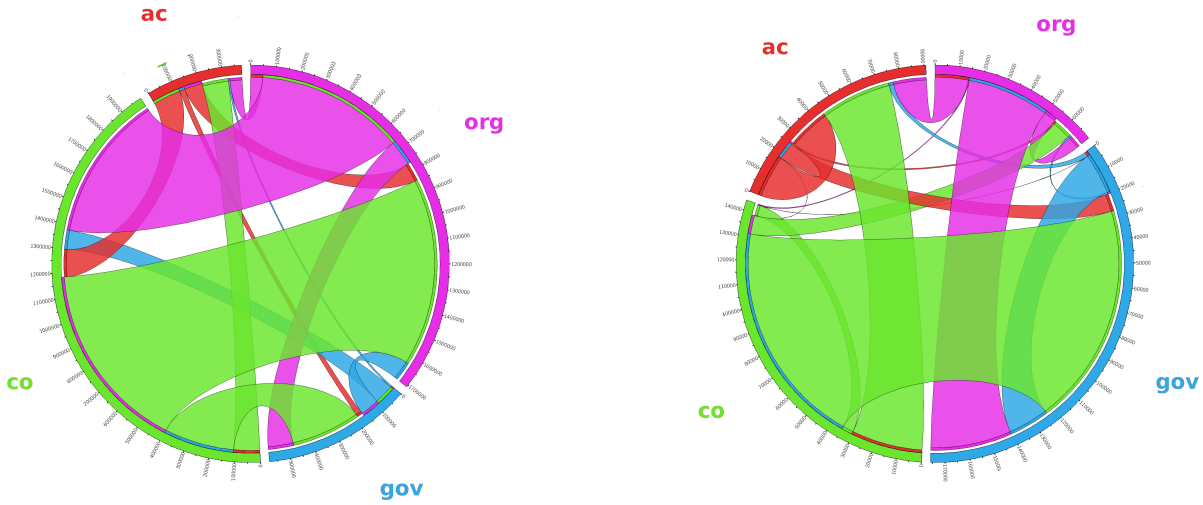

(a) (b)

Figure 4: Links between four second-level domains. Panel *a* shows the absolute number of links between different SLDs (self-loops are excluded), and panel *b* shows the relative number of links normalized by the size of target subdomain.

tween universities' linking practices and three variables: institutional affiliation, league table ranking, and geographic location. Our hypothesis in doing so was that higher status academic institutions would be more strongly linked to than lower status institutions and would also be more strongly interconnected with their peer institutions.

For the analysis, we built a list of the 121 universities listed in the most recent Sunday Times University Guide.[8] Each of these universities has a website, all of which use the `.ac.uk` suffix. We obtained the third-level domain (e.g., `ox.ac.uk`) for each. Further data collection as necessary is described in the respective subsections that follow.

4.3.1 *Group affiliation*

Many British universities belong to associations, formed to represent their interests and facilitate collaboration. The groups are neither mutually exclusive nor exhaustive, meaning that universities can belong to none, one, or more than one group, but for practical and political reasons most universities belong to only one. We collected data on the memberships of five groups, the Russell Group,[9] the 1994 Group,[10] the University Alliance,[11] the Million+ Group,[12] and the Cathedrals Group.[13]

The best known of these is perhaps the Russell Group, formed in 1994 and now constituted of 24 members. The 1994 Group, which represented smaller research institutions,

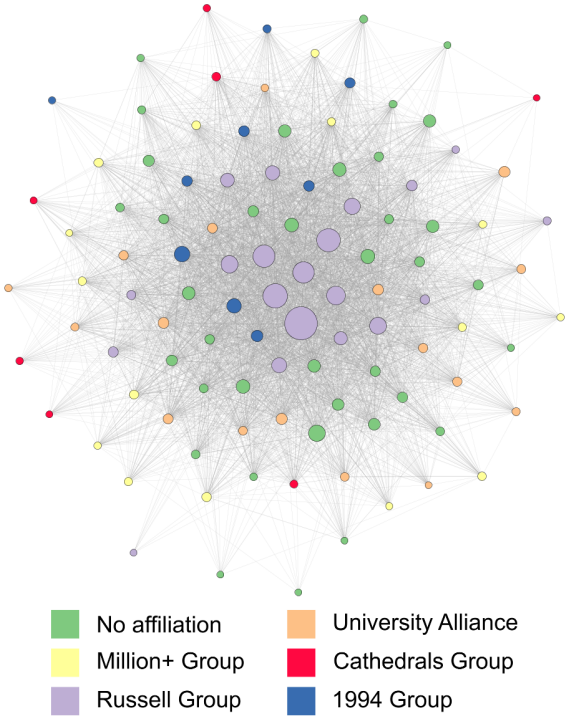

Figure 5: Network diagram of hyperlinks between universities. Different colors indicate different university affiliations.

[8] `http://www.thesundaytimes.co.uk/sto/University_Guide/`
[9] `http://www.russellgroup.ac.uk/our-universities/`
[10] `http://www.timeshighereducation.co.uk/news/was-1994-groups-demise-triggered-by-relaunch-delays/2008999.article`
[11] `http://www.unialliance.ac.uk/member/`
[12] `http://www.millionplus.ac.uk/who-we-are/our-affiliates/`
[13] `http://cathedralsgroup.org.uk/Members.aspx`

was formed in response to the Russell Group, but disbanded in 2013; given the time frame of the dataset we include the 11 final members of the group in our analysis. Of the final three groups, the University Alliance is formed of 22 business-oriented UK universities, the Million+ Group is made up of 17 mostly 'new' (post-1992) institutions, and the Cathedrals Group is made up of 16 universities originally instituted as church-led teacher training colleges. The stated purposes of these groups differ somewhat, but each are constituted broadly to serve the research interests of their members.

In comparing group membership to the density of links between different universities, we sought to discover whether academic affiliation was associated with the density of links between institutions. To do this, we performed a network analysis, investigating whether the universities clustered on the basis of group affiliation. Figure 5 shows a network diagram, with different affiliations marked by different colors.

To the naked eye, Figure 5 shows no discernible clustering on the basis of group affiliation, and network analysis bears this out. The division of the network by affiliations has a modularity score [16] of -0.003, indicating that the division of the network into clusters based on university affiliation is no better than dividing the network into five random clusters. On an individual basis, only one group, the Russell Group, has many internal links and comparatively fewer links to institutions outside the group. It is the most strongly connected group with an internal hyperlink density of 0.71. The Russell Group, which includes 24 of the leading international UK universities with some of the highest levels of research funding, arguably represents most if not all of the elite universities in the UK. It contains 9 of the 10 top-ranked UK universities, including both Oxford and Cambridge. That these universities are more strongly linked to each other is likely related at least in part to their active research cultures, with many collaborations between researchers at these top institutions. The lack of strong web connections in the other associations, however, suggests that while these institutions may or may not have strong connections among their members by other measures, there is no evidence that universities strongly link to the websites of institutions with which they share group affiliation over other institutions.

4.3.2 *League table ranking*

University league tables are an important if imperfect indicator of a university's prominence. Modern league tables incorporate a whole range of measures, including factors related to teaching, research, and student satisfaction. As such, we wanted to investigate whether a university's league table ranking is associated with its web presence, and whether the relationship has changed over time, in terms of both increasing adoption and development of an institution's web presence and its changes in league table ranking over time. For this analysis, we collected the rankings of British universities published in The Times Good University Guide for three years, 2000, 2005 and 2010, and compared these rankings with data from crawls conducted in the same three years.

In conducting the analysis, we used ten common measures of network centrality for each of the three different years to gauge the relationship between each university's league ranking and its position in the network of hyperlinks flowing between university third-level domains. We then produced

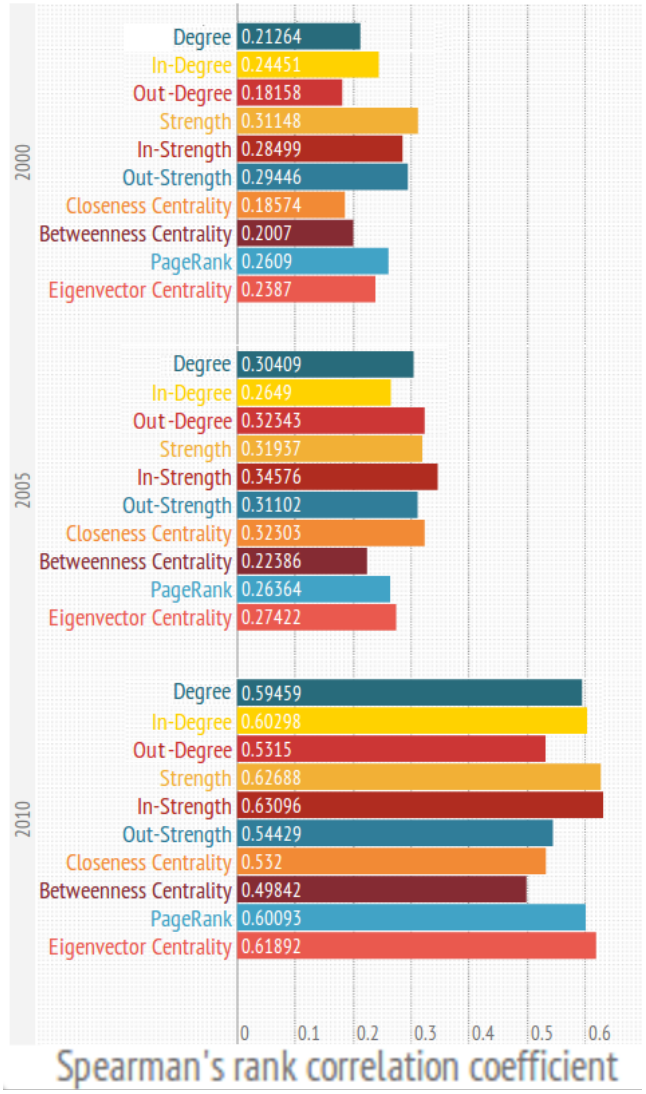

Figure 6: Spearman's rank correlation coefficients between university league table rankings and ten different network centrality measures for three years.

lists ranking the universities for each year by each centrality measure and computed Spearman's rank correlation coefficient for each centrality ranking and league table ranking combination. These correlation coefficients are shown in Figure 6.

For most measures of centrality used, a pattern emerges: the data for 2010 shows the strongest correlation between league table ranking and centrality, while the relationship is less evident for 2000 and 2005. The most strongly correlated correlation measure is in-strength, a sum of all the hyperlinks linking to a given web domain. This measure uses the weight of each edge, which corresponds to the number of hyperlinks between any two third-level domains. This differs from in-degree which measures the number of other domains that link to a given web domain. Figure 7 shows the fairly strong correlation between universities' league table rankings and their network positions as measured by in-strength.

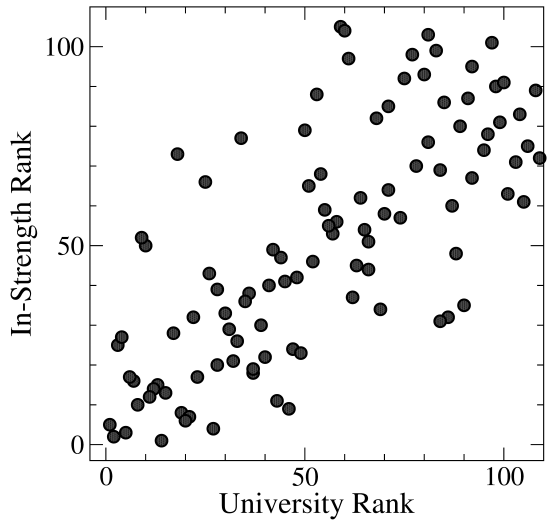

Figure 7: University in-strength rankings compared to university league table rankings for 2010. Spearman's rank correlation is 0.63.

What Figure 6 and Figure 7 suggest is two-fold: first, that university prominence, as measured by league table position, is an increasingly stronger predictor of the number of links to that institution over the 2000–2010 period. Whether this is an example of the Matthew Effect ("the rich get richer") [15] whereby highly prominent institutions become well-linked institutions largely as a result of their prominence (and conversely, marginal institutions become more marginalized as a result of their lack of prominence), or whether there is another independent factor at play here cannot be determined from these data. However, the second conclusion is clear: the hyperlink patterns within the UK academic subdomain support the notion that the web does not inherently challenge existing power structures. Instead, the saturation of the `.ac.uk` subdomain, in terms of the presence of essentially all possible academic institutions by 2003 (Figure 1), has resulted by 2010 in a subdomain in which network centrality closely mirrors prominence as measured by league tables.

4.3.3 Role of geography

Finally, we investigated whether any association exists between the geographic proximity of British universities and the density of hyperlinks between them. This analysis builds upon work by Pan et al. [17] who found, at a global scale, that rates of academic citations and collaborations between two cities diminish as the distance between them increases, following gravity laws. We conduct a similar analysis, replacing citations and collaborations with hyperlinks collected in the web domain data.

We collected geographic coordinates for each of the British universities in the list using a simple Google Maps search. Universities can be spatially complex, sometimes having multiple campuses and satellite sites; so, some discretion was occasionally required in identifying the center of each university.

The standard, naïve gravity law approach would suggest that the number of hyperlinks, or the strength of the connection, between two given universities is inversely proportional to the square of the distance between the two universities. We let S_{ij} denote the strength from university i to university j. Focusing on the data from 2010, the left frame of Figure 8, shows that the relationship between this measure and the geographical distance between the two universities is very noisy. To correct for the different sizes of universities and their different linking practices (some universities may just link more than others), we normalize these strengths. We divide S_{ij} by the sum of the weights of all edges coming from university i (S_i^{out}) multiplied by the sum of the weights of all edges linking to university j (S_j^{in}). We denote this normalized measure σ_{ij} and plot it against physical distance in the right frame of Figure 8. With this normalization, the relationship between distance the number of hyperlinks (strength) between universities is very clear. In both frames, we use a moving average window with a length of 500 data points and therefore a lower bound of 20km is introduced. An upper bound is induced by considering only the universities within the UK in this study. However, the gravity law holds significantly within a distance range of almost two decades.

Letting d_{ij} denote the geographical distance between two universities, we then seek the exponent a, which best fits the observed data following $\sigma_{ij} \propto d_{ij}^{-a}$. Using the least squares methods, we fit a linear function to the logarithmically transformed data and find $a = 0.28 \pm 0.02$, which closely matches the findings of Pan et al. [17] for citation and collaboration networks. In that study, Pan et al. found an exponent of $a = 0.30$ for the citation network before any normalization, while finding an even stronger role for geographical distance ($a = 0.77$) after applying a similar normalization to the one we apply here.

Figure 9 maps the universities in the sample along with the connections between them colored according to σ. It is evident, specially in the map of 2010, that the longer connections generally have weaker strength. It is worth nothing that the size limit of the dataset and the geographical constraints, such as the dense region of London extended to Oxford and Cambridge, which includes a large number of universities in our dataset, could partially drive the strong geographical dependency we observed. This dense region is particularly visible in the map of 2005 in Figure 9.

5. CONCLUSION

In this paper, we have reported some of the first findings based on longitudinal analysis of the entire recorded history of the UK web domain. While this current analysis is by necessity at a macro-level in terms of detail, it nevertheless demonstrates the potential of these data for detecting changes in patterns in web linking behavior over time, evidence related to the growth and expansion of the web, and uneven patterns of linking within subdomains, such as the academic `.ac.uk` subdomain discussed here. We have shown that even though the growth of the commercial side of the web has resulted in increasing commercial dominance of the UK web space in terms of absolute number of nodes, the academic and government subdomains receive proportionally more inlinks per domain. In examining the academic subdomain in particular, we have shown that while there is no generalized clustering based on the affiliation of academic institutions, there are clear patterns in terms of higher inlinks to the highest status academic institutions and stronger connections between geographically-closer institutions.

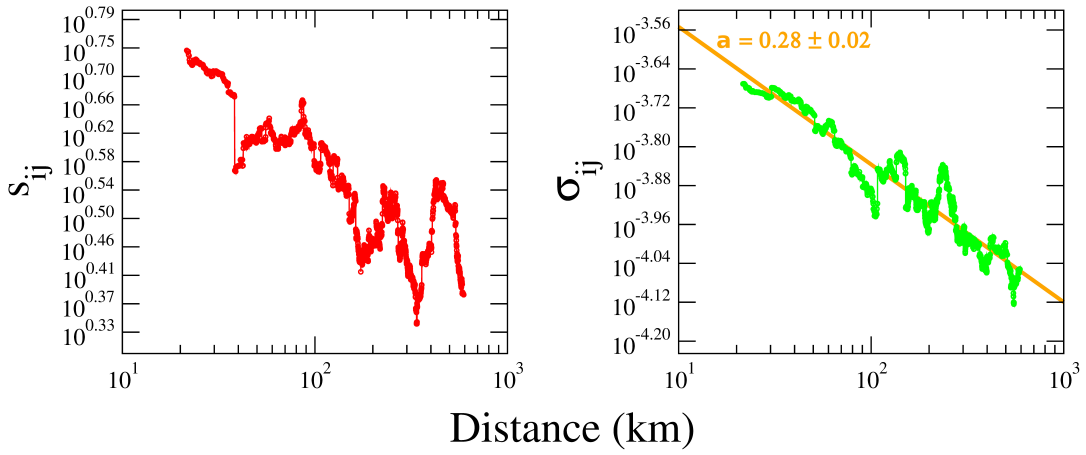

Figure 8: Left: Raw hyperlink strength (S_{ij}) between universities versus geographical distance, and Right: Normalized hyperlink strength ($\sigma_{ij} = \frac{S_{ij}}{S_i^{out} S_j^{in}}$) between universities versus geographical distance. The normalized measure follows a gravity-law model with an exponent of $a = 0.28 \pm 0.02$.

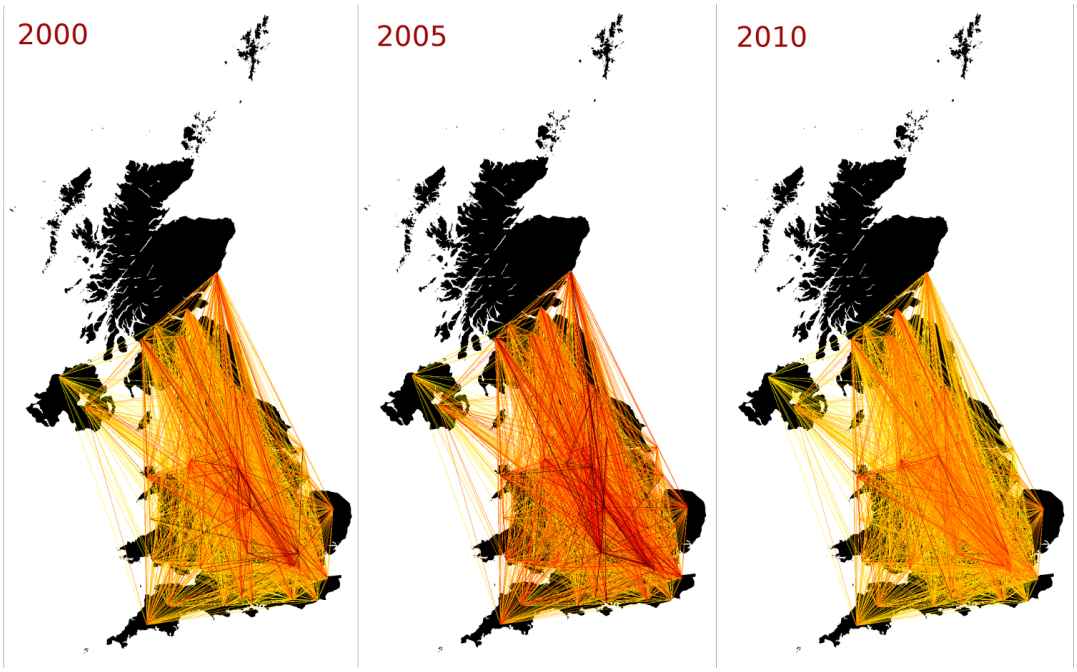

Figure 9: Maps of the UK universities under study for three years: 2000, 2005, and 2010. The connections are the hyperlinks, and color corresponds to the normalized strength of each link (σ_{ij}). The reddest links correspond to the strongest connections.

This analysis also suggests many future possibilities for research with these web archive data, including more detailed micro-level analysis of linking behavior within various subdomains over time, discovery of networks of collaboration between subunits of institutions, comparison between link measures and other measures of prominence such as citation networks, and analysis of other subdomains besides `.ac.uk`. In addition, there are ongoing efforts to prepare the full-text corpus extracted from the web archive for research (rather than the link corpus used here), which will be able to be combined with these data to answer more detailed questions about the content of the web, the context for links, and discourses on the web.

6. ACKNOWLEDGMENTS

The authors would like to thank Ning Wang for his advice and support on data cleaning and Andreas Kaltenbrunner for his help with creating the geographic visualizations. The authors are also grateful for funding from UK Jisc for the "Big Data: Demonstrating the Value of the UK Web Domain Dataset for Social Science Research" grant (16/11 Enhancing the Sustainability of Digital Collections) that supported the data extraction and early analysis, and further funding for analysis from the UK Arts and Humanities Research Council for the "Big UK Domain Data for the Arts and Humanities (BUDDAH)" grant (AH/L009854/1). Finally, the authors would like to thank our anonymous reviewers for their helpful comments on an earlier version of this paper.

7. REFERENCES

[1] S. G. Ainsworth, A. Alsum, H. SalahEldeen, M. C. Weigle, and M. L. Nelson. How much of the web is archived? In *Proceedings of the 11th Annual International ACM/IEEE Joint Conference on Digital Libraries*, pages 133–136. ACM, 2011.

[2] R. Baeza-Yates, C. Castillo, and E. N. Efthimiadis. Characterization of national web domains. *ACM Transactions on Internet Technology (TOIT)*, 7(2):9, 2007.

[3] I. Bordino, P. Boldi, D. Donato, M. Santini, and S. Vigna. Temporal evolution of the UK web. In *Data Mining Workshops, 2008. ICDMW '08. IEEE International Conference on*, pages 909–918, Dec 2008.

[4] N. Brügger. Historical network analysis of the Web. *Social Science Computer Review*, 31(3):306–321, 2013.

[5] N. Brügger. Probing a nation's web sphere: A new approach to web history and a new kind of historical source. 2014.

[6] M. Dougherty and E. T. Meyer. Community, tools, and practices in web archiving: The state of the art in relation to social science and humanities research needs. *Journal of the American Society of Information Science & Technology*, 2014.

[7] M. Dougherty, E. T. Meyer, C. Madsen, C. V. den Heuvel, A. Thomas, and S. Wyatt. Researcher engagement with web archives: State of the art. Technical report, 2010.

[8] T. Escher, H. Margetts, V. Petricek, and I. Cox. Governing from the centre? Comparing the nodality of digital governments. In *Annual Meeting of the American Political Science Association*, 2006.

[9] K. A. Foot and S. M. Schneider. *Web Campaigning*. The MIT Press, 2006.

[10] S. A. Hale. Net increase? Cross-lingual linking in the blogosphere. *Journal of Computer-Mediated Communication*, 17(2):135–151, 2012.

[11] S. A. Hale, T. Yasseri, and H. Margetts. Extracting clean hyperlink and website data from the JISC UK Web Domain Dataset. Technical report, 2014.

[12] B. Kahle. Preserving the Internet. *Scientific American*, 276(3):82–83, 1997.

[13] B. M. Leiner, V. G. Cerf, D. D. Clark, R. E. Kahn, L. Kleinrock, D. C. Lynch, J. Postel, L. G. Roberts, and S. Wolff. A brief history of the Internet. *ACM SIGCOMM Computer Communication Review*, 39(5):22–31, 2009.

[14] J. Masanès. Web archiving: Issues and methods. In J. Masanès, editor, *Web Archiving*, pages 1–54. Springer, 2006.

[15] R. K. Merton. The Matthew effect in science. *Science*, 159(3810):56–63, 1968.

[16] M. E. J. Newman. Modularity and community structure in networks. *Proceedings of the National Academy of Sciences*, 103(23):8577–8582, 2006.

[17] R. K. Pan, K. Kaski, and S. Fortunato. World citation and collaboration networks: Uncovering the role of geography in science. *Scientific Reports*, 2, 2012.

[18] N. Payne and M. Thelwall. Longitudinal trends in academic Web links. *Journal of Information Science*, 34(1):3–14, 2008.

[19] R. Rogers and N. Marres. Landscaping climate change: A mapping technique for understanding science and technology debates on the World Wide Web. *Public Understanding of Science*, 9(2):141–163, 2000.

[20] M. Thelwall, R. Tang, and L. Price. Linguistic patterns of academic Web use in Western Europe. *Scientometrics*, 56(3):417–432, 2003.